

DirAc: An Integrated Circuit for Direct Acquisition of the M-Code Signal

John W. Betz, John D. Fite, and Paul T. Capozza
The MITRE Corporation

BIOGRAPHIES

John W. Betz is a Fellow of The MITRE Corporation. He received a Ph.D. in Electrical and Computer Engineering from Northeastern University. He led the GPS Modernization Signal Design Team's Modulation and Acquisition Design Subteam, contributing to many aspects of M-code signal design and evaluation, including developing the binary offset carrier modulation for the GPS M-code signal. He has continued work on many aspects of modernized signal designs, and has recently been involved in technical aspects of the negotiations between US and EC on GPS and Galileo. He was co-recipient of the ION Burka award in 2001.

John D. Fite received a B.A. degree in Mathematics and Mathematical Sciences in 1980 and a Masters of Electrical Engineering in 1981 from Rice University, Houston, Texas. Mr. Fite is a Principal Engineer at The MITRE Corporation, McLean, VA, where his work has included the design and development of rapid acquisition algorithms for GPS signals, multi-user detection algorithms, and adaptive array processing techniques.

Paul T. Capozza is the Director of the Microelectronics Technical Center at The MITRE Corporation. Mr. Capozza is responsible for leading the design and development of complex integrated circuits for communications, navigation, and radar programs. Hardware developed for these programs includes Systems-on-a-Chip, RF and microwave design, and optics. Recently for the GPS JPO, he defined next generation integrated circuit architectures for the M-code signal receivers and he has developed GPS anti-jam integrated circuits.

ABSTRACT

This paper describes the first integrated circuit (IC) designed, fabricated, and tested to perform direct acquisition of the M code signal. This DirAc IC prototype provides direct acquisition capability for test receivers and also demonstrates the feasibility of performing direct acquisition over extended regions of time and frequency uncertainty. The IC is designed and fabricated using 180 nm technology, and has been tested to demonstrate complete functionality and full performance. It uses parallel code matched filters, with FFT-based backend

processing to search over 800 Hz of frequency uncertainty and 10 msec of time uncertainty in parallel, using off-chip memory for noncoherent integration. Multiple such time-frequency tiles are searched serially. Inputs are sampled at 2 bits each inphase and quadrature. The DirAc IC supports a maximum integration time (combined coherent and noncoherent integration) of 1.28 seconds, and includes compensation for code Doppler. Coherent integration time up to 10 msec can be used. The DirAc IC's architecture takes advantage of the M-code signal's binary offset carrier (BOC) modulation to reduce acquisition processing complexity. DirAc supports different modes and features of the M-code signal. Hardware is time-shared between inphase and quadrature processing and also between upper and lower sidebands of the BOC modulation. The architecture uses a pipelined design to provide the equivalent processing capability of 3.2 million parallel correlators, performing 2 tera operations per second. Average power consumption in a typical application is less than 1 mW. The IC design and layout process are also described, identifying techniques used to efficiently design and layout the IC. Theoretical predictions are provided for search speed and for the ability to work at different levels of carrier-to-noise density ratio.

INTRODUCTION

The M-code signal, the modernized GPS military signal designed in the late 1990s, is scheduled to be first transmitted by a Block IIR-M satellite in 2005. As described in [1], the M-code signal's revolutionary design includes a novel modulation [2], new data message, and new security architecture. M-code signal acquisition relies primarily upon direct acquisition, where in effect the receiver correlates (over time and frequency shifts) a locally generated replica of an M-code signal with the received waveform. When there is a match between the replica and a received signal, coarse synchronization is achieved, and the receiver commences signal tracking, data message demodulation, and position calculation. Since the M-code signal, like the current GPS military signal called Y-code signal, uses very long spreading codes, signal acquisition cannot take advantage of the short spreading codes that simplify acquisition processing in civilian signals such as the GPS C/A-code signal.

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2004		2. REPORT TYPE		3. DATES COVERED 00-00-2004 to 00-00-2004	
4. TITLE AND SUBTITLE DirAc: An Integrated Circuit for Direct Acquisition of the M-Code Signal				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) MITRE Corporation, 202 Burlington Road, Bedford, MA, 01730-1420				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 10	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

While direct acquisition circuits were developed for Y code receivers in the 1990s [3], these circuits provided much less capability than would be needed for direct acquisition to be the primary mechanism for acquiring the M-code signal. During the design of the M-code signal, studies demonstrated that a combination of factors would allow direct acquisition to surpass the design requirements for the M-code signal. But the results of these studies did not lead to consensus that ICs ready for receiver production in the latter half of this decade could meet the performance requirements while providing adequately low complexity, low parts cost, low peak and average power consumption, and low thermal dissipation. Functioning silicon was needed to remove remaining doubts.

A team with expertise in systems engineering, digital signal processing, and IC design took on the challenge of developing a prototype IC for direct acquisition of the M-code signal. The team identified ways to exploit the unique characteristics of the M-code signal, evaluated processing architectures that balanced risk and capability, developed predictions of performance and of IC complexity, designed and applied algorithms, developed detailed simulations, and traded off processing implementations, yielding design files that were sent to the foundry only 12 calendar months after the design effort began. The resulting “DirAc” ICs have been packaged and tested, confirming that they provide full functionality and meet or exceed performance predictions. Software and hardware development is underway to integrate the IC into a test receiver for further testing.

The next section of this paper discusses direct acquisition of the M-code signal, outlining issues and opportunities to be considered. The following section describes the DirAc architecture. The subsequent section describes the first DirAc IC’s design, including digital signal processing, architecture, and layout, built using 180 nm lithography readily accessible in 2001. The succeeding section outlines a second version DirAc that could be developed using 130 nm technology available in 2003. Fundamental performance characteristics are provided in the subsequent section, while the final section summarizes the findings of this paper.

DIRECT ACQUISITION OF THE M-CODE SIGNAL

Signal acquisition involves the steps that take a receiver from a state of being powered on and having passed self-test, to providing an initial estimate of position, time, or velocity (PVT) at specified accuracy. Time to first fix (TTFF) denotes the delay between starting the acquisition process and providing PVT with specified accuracy.

In conventional receivers, TTFF then involves the time for coarse initial synchronization (obtaining initial

alignment between the receiver’s timing and frequency and those of the received signal), signal tracking or other processing that produces refined and repeated estimates of signal timing and frequency, reading the data message to obtain position and time at the satellite transmitting the signal (if needed), obtaining signal tracking and satellite position and time for three or more additional signals, and then calculating PVT using the estimates of signal timing and frequency along with positions and times at the satellites.

Even though coarse initial synchronization is only one part of acquisition processing, this paper complies with common terminology, calling the circuit that performs coarse initial synchronization an “acquisition circuit.”

Typically, an acquisition circuit crosscorrelates a locally generated signal replica against the received waveform containing multiple signals, interference and possibly jamming, and noise. In concept, the locally generated reference is shifted in time and frequency, forming a segment of a cross ambiguity function (CAF) [4] between the replica and the desired received signal. The time duration of the signal segments used in the crosscorrelation is called the coherent integration time. Noncoherent integration can be accomplished by adding the magnitudes of multiple CAFs, computed over the same ITU and IFU. This noncoherent integration enhances performance in noise and jamming, but consumes additional time to collect and process the longer segment of received waveform.

Digital processing actually searches discrete values in time and frequency space, called time-frequency cells. The time span and frequency span searched in parallel by an acquisition circuit may be called a time-frequency tile—composed of multiple cells in a rectangular array. If the ITU or IFU is larger than the span of the tile, sequential tiles are computed serially to compute the CAF over the entire ITU and IFU. Figure 1 shows how cells and tiles fit into the ITU and IFU. Although for signals with short periodic spreading sequences, the largest ITU that must be searched corresponds to the period of the spreading sequence, signals whose spreading sequences have much longer periods require search of the entire ITU.




Figure 1. Time-Frequency Tiles Filled with Cells Are Used to Search a Region of Initial Time Uncertainty and Initial Frequency Uncertainty

The magnitude CAF, computed with or without noncoherent integration, is used to form a test statistic. A threshold setting algorithm establishes a criterion, and the time and frequency of any magnitude-squared CAF value that exceeds this threshold criterion correspond to a possible coarse initial synchronization. Some of these detection reports may be false, since establishing the threshold value too high excessively reduces the probability of a valid detection. Thus, some of the reported detections correspond to false initial synchronization points, and various techniques are used to distinguish between valid and false synchronization points.

If spacing between time-frequency cells is too wide, a peak in the CAF signifying coarse initial synchronization may occur in between sample points, degrading the opportunity to detect this peak. In general, modulations having narrower peaks must be sampled faster to avoid this problem in the time domain. Sampling in the frequency domain is independent of modulation design, but instead is proportional to the reciprocal of the coherent integration time used in computing CAFs—longer coherent integration times require finer frequency spacing.

When custom hardware is used for size and power efficiency, crosscorrelations are often implemented in the time domain, since two or fewer bits of quantization are needed, and less silicon area is needed. When the crosscorrelations used to compute the CAF are implemented using time-domain computations, the number of arithmetic operations required to compute a CAF that covers a given ITU increases with the square of the sampling rate. Thus, for signals whose modulation is binary phase shift keying with rectangular spreading symbols, the computational burden (quantified as the rate of arithmetic operations) increases with the square of the spreading code rate. Since the correlation peak is very narrow for signals with binary offset carrier (BOC)

modulation, the same logic would motivate even higher sampling rates, and thus higher computational burdens, particularly when the subcarrier frequency is greater than the spreading symbol rate.

Fortunately, sideband acquisition processing of BOC modulations [5] significantly reduces the computational burden for signals having BOC modulations with subcarrier frequency greater than the spreading code rate. Since the signal spectrum has distinct upper and lower sidebands, they can be separately filtered, downconverted to DC, and decimated at a sample rate commensurate with the spreading code rate, independent of the subcarrier frequency. Separate CAFs can be formed from the resulting waveforms from the upper and lower sidebands, using as a replica signal the spreading sequence without subcarrier modulation. The CAFs from upper and lower sidebands are noncoherently integrated.

Figure 2 shows a conceptual block diagram of this processing for the M-code signal. The computational advantages from processing for coarse initial synchronization are significant. For M-code signal, sideband acquisition processing requires approximately 4% of the computational load required by wideband acquisition processing of the M code signal, and 50% of the computational load required by acquisition processing of the Y-code signal using the same coherent integration time and the same number of noncoherent integrations to search the same ITU and IFU. Storage requirements for sideband acquisition processing of the M-code signal are also significantly less than those for either wideband acquisition processing of the M-code signal, or acquisition processing of the Y-code signal.

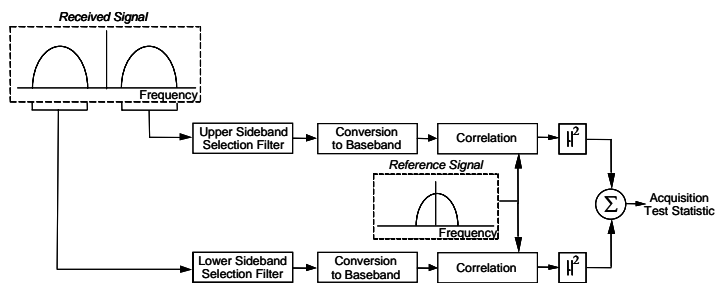


Figure 2. Sideband Acquisition Processing of the M-Code Signal

DIRAC ARCHITECTURE

This section describes the DirAc architecture. An overview is provided first, followed by more detailed discussion of separate portions.

DirAc Architecture Overview

The DirAc signal processing architecture is shown in Figure 1. The input is four separate sampled input streams: the sampled inphase upper sideband (I_{USB}), the

quadrature upper sideband (Q_{USB}), the inphase (I) lower sideband (I_{LSB}), and the quadrature (Q) lower sideband (Q_{LSB}), each clocked at 5.115 MHz. These samples streams are converted to a single interleaved stream at 20.46 MHz in order to reuse the code-matched filter (CMF) hardware.

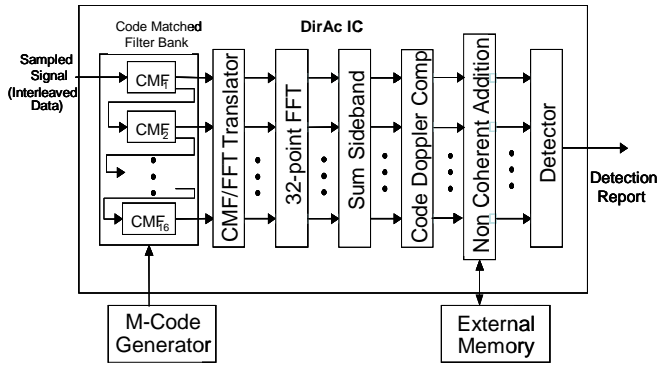


Figure 3. DirAc Architecture for Rapid Acquisition Using a Bank of Code Matched Filters

The CMF bank consists of 16 short-time CMFs, each providing 0.625 msec coherent integration time, that crosscorrelate the received input samples with the reference M-code samples. Input samples from one CMF are shifted into the next CMF at the sample rate. Each CMF computes a short-time crosscorrelation that is sent to the CMF/FFT translator. The CMF/FFT translator converts the interleaved I_{USB} , Q_{USB} , I_{LSB} , Q_{LSB} crosscorrelations from each CMF into parallel I, Q pairs of USB and LSB data which are sent to the FFT. The translator processes each of the CMF outputs in parallel, producing a total of sixteen parallel channels.

The fast Fourier transform (FFT) performs a 32-point zero-padded complex FFT of the I and Q data. Zero-padding interpolates between frequency bins to reduce scalloping loss caused when the true frequency value falls between the discrete frequencies computed in the CAF. The FFT provides a coherent integration time of $16 \times 0.625 = 10$ msec while also computing the CAF over different frequency values. Each frequency bin in the FFT produces data in 100 Hz frequency bins, with adjacent bins overlapped by 50 Hz. To further reduce losses [6], only the 16 center bins in the FFT are retained, covering ± 400 Hz.

Magnitudes of the separate CAFs computed from upper and lower sidebands are then summed, forming a single magnitude CAF over 10 msec and ± 400 Hz.

When the noncoherent integration time is long, code Doppler can degrade performance unless compensation is used. Code Doppler compensation (CDC) prevents time smearing of a CAF peak due to differences between the receiver and satellite spreading code rates.

Noncoherent integration improves the detection performance by summing magnitude CAFs computed to search the same time-frequency tile using successive 10 msec segments of received waveform, reducing the variance of the resulting CAF. The addition process requires external storage of the intermediate time-frequency tiles.

After the requisite number of noncoherent integrations, detection processing examines the resulting magnitude CAF to determine if there are candidate coarse initial synchronizations in the magnitude CAF. A detection threshold is established; and for each magnitude CAF values exceeding that threshold, a detection report is provided containing the magnitude value and the time and frequency values of the cell. If adjacent cells also exceed the threshold (a common occurrence when the signal energy straddles multiple cells), the information pertaining to adjacent cells is also included in the detection report.

CMF Bank Description

Simpler versions of the CMF bank design can be found in [6, 7]. An advantage of the CMF bank design implemented in DirAc, as compared to other parallel cross-correlation methods (e.g. [8, 9, 10]) is that it removes the need for intermediate storage of input samples and partial coherent correlation sums often found in other designs. When all the multiply-accumulate operations needed for the coherent integration are implemented in hardware, there is no need for intermediate storage to cache the partial cross-correlation sums. This is in contrast to architectures that use parallel active correlators or the hybrid active-passive cross-correlation methods, which reuse arithmetic resources for computing the coherent integration but require intermediate memory storage. Since all the needed partial cross-correlation values are provided to the FFT at the same time, there is no need for the intermediate memory to store partial correlation sums.

The tap structure in the short-time CMFs, shown in Figure 4, exploits the fact that multiple input data streams (I_{USB} , Q_{USB} , I_{LSB} , Q_{LSB}) are being crosscorrelated with the same reference spreading sequence. Interleaving the data and running at 20.46 MHz allows the same multiplier, code registers, and adder tree to be reused for each input data stream. Each short-time CMF has 3197 taps, and each tap structure includes a shift register of length 4, a multiplier, and two code registers. The two code registers allow for a seamless transition from one code to another. This is especially advantageous when noncoherent integration over multiple time segments is required.

Each CMF consists of multiple tap structures connected to a dual adder tree structure. One adder tree sums up the products corresponding to the odd numbered taps and the

other to the products corresponding to the even numbered taps. The resulting sums are kept separate and can be noncoherently added together after the FFT processing to support M-code signal characteristics [11].

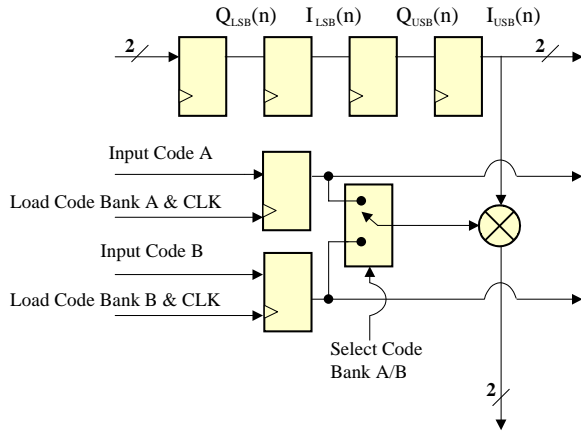


Figure 4. Tap Structure Using Hardware Sharing across I,Q Samples for USB and LSB

Noncoherent Integration over Time

After the FFT processing that computes separate CAFs for upper and lower sidebands, the magnitude CAFs are summed together. The CAFs corresponding to the odd and even taps are also summed together noncoherently.

In the current design, every 10 milliseconds the CMF bank and FFT processing generates a 51150 x 16 sample corresponding to 10 msec by 800 Hz time-frequency space. If there is an ITU greater than 10 msec and noncoherent integration is not needed, then the same code reference PN sequence remains in the CMF taps for the entire search time. However, if noncoherent integration over multiple 10 msec intervals is required, the reference PN code is updated at every 10 msec interval, so the code offset for the current tile is the same as the previous tile.

This noncoherent integration occurs after code Doppler compensation. The multiple magnitude CAFs corresponding to the same code offsets but computed at different times are summed together before detection. The first magnitude CAF is stored on the off-chip memory and then, for subsequent integrations, the stored magnitude CAF is added to the current magnitude CAF and then the sum is written back into the off-chip memory.

Code Doppler Compensation Function

Prior to summing the current and stored magnitude CAFs, some frequency-dependent delay compensation is required. Any frequency offset between transmitter and receiver results in a time compression or expansion (known as companding). This companding affects both the carrier frequency and the spreading code rate of the received signal. While offsets in carrier frequency are addressed by the FFT, the change in spreading code rate

produces a lack of correlation between the local PN code and the received signal over long integration times.

The loss of coherence due to time companding of the baseband signal can be compensated by the use of short-time correlations followed by post-processing. (e.g. [12, 13]). The companding effect causes a correlation peak to “drift” along the time axis of the short-time correlation tile at different times. If no compensation is made for the drift of the correlation peak as a function of code Doppler offset and as the number of noncoherent integrations, increasing the number of integrations beyond a certain point provides no additional benefit. Code Doppler compensation predicts the relative location of the correlation peak from correlation block-to-block and to apply the necessary delays to make sure the correlation peaks remained aligned from tile to tile.

Code Doppler compensation consists of a bank of integer and fractional delay lines that are used to maintain the correct peak alignment from tile to tile. The fractional delay lines employ a 4-tap Lagrangian interpolator that uses a table to assist in proper delay coefficient selection. As each CAF is processed, the integer delay lines and fractional delay lines are initialized and/or updated to counteract the correlation peak drift. The delay of the integer variable delay line is a function of both the code Doppler offset and the number of integrations performed.

DIRAC IC DESIGN DETAILS

VLSI implementation decisions for the DirAc IC were driven by the CMF bank since it represents 95% of the hardware resources and processing. This section provides an overview of the IC architecture and discusses the important issues of clock and power distribution.

DirAc IC Architecture Overview

Two implementation strategies were evaluated, hardware reuse and massive parallel processing. An architecture implementing a hardware reuse strategy stores and processes in an iterative fashion using a subset of processing elements in this case correlator taps at an increased rate. This approach limits the required hardware for processing but requires overhead for coordinating the reuse. The major VLSI implementation drawbacks with this approach, especially when implementing a CMF bank, is the requirement to store input samples and intermediate partial correlations while reusing the hardware. In addition, high-speed clock management is required to obtain the reuse factor. Memory bandwidth requirements also force the architecture to support multiple memories on chip. These high-speed memories are challenging to place and route due to route congestion, global routes that introduce signal integrity issues, and routing blockages.

The alternative architecture in contrast facilitates VLSI implementation. We selected a parallel datapath

implementation of the CMF. Figure 5 depicts the DirAc IC layout with an overlay of the 16 CMFs and their signal data flow. In this architecture the CMF is processed in parallel. The systolic processing associated with this architecture minimizes global routing and in fact is by definition limited to local routing. This has two benefits: it eliminates global routes minimizing signal integrity issues, and it minimizes power dissipation related to interconnect. In addition, the inherent symmetry of the systolic array facilitates clock and power distribution. Another major benefit is that since the 16 CMF elements are identical the CMF designs can be reused, simplifying the design specification, layout, and verification process allowing for a timely implementation.




Figure 5. DirAc IC's Systolic Array Architecture

Clock Distribution

The DirAc IC requires 519,500 register elements, representing 43% of the die's core logic area as shown in Figure 6 and over 60% of the total number of transistors.




Figure 6. DirAc Layout Highlighting Die Area Associated with Registers

In order to minimize clock skew without impacting overall design performance, clock distribution was addressed at two levels: macro-level clock distribution and top-level. Clocks were routed on the top two metal layers for each macro (i.e., a CMF). Restricting clock routing to metal 5 and 6 opened routing resources at the lower metal layers for signal routing. All clock routes were also double spaced to minimize interference to signal routes. State-of-the-art place and route tools automatically introduced a clock tree for each macro.

At the top level, clock trees were hand placed. Then an automated router was used to distribute clock interconnects. These routes were double spaced and interleaved with ground.

The final clock network has a clock skew of only 83 psec. The clock tree contains 10 levels and over 15,000 buffers. Power dissipation related to the clock tree was on the order of 10% of total dynamic power dissipation.

Power Distribution

A major design consideration for the DirAc IC was distributing power across the die. The DirAc IC's IR drop requirement had to be limited to 25 mV. Failure to meet this requirement would have potentially resulted in timing errors or logic failures.

The power grid was constructed based on three design criteria: achieving an IR drop below 25 mV, balancing power grid routing tracks with signal routing requirements, and not interfering with the placement and routing of the clock tree.

In order to gain insight early on in the power grid design process a power grid analysis tool was developed using SPICE. Estimates of macro power dissipation based on RTL power analysis were incorporated with a resistor and current source network that modeled the power routes and power dissipation. Insight from this tool drove the final macro placement, course and fine power grid size, and their connections. Multiple iterations and modifications to the power grid resulted along with architectural modifications to the DirAc IC.

In the end, the SPICE-based analysis tool predicted 18 mV final IR drop. We confirmed this IR drop using a commercial gate-level power analysis tool after the DirAc IC was completely placed and routed. Figure 7 shows an IR drop map for the DirAc IC. The largest IR drop occurs at 22 mV. Note the symmetric shape of the IR drop and the hot spot is located towards the center of the die.




Figure 7. DirAc IC IR Drop Map

DIRAC CHARACTERISTICS AND PERFORMANCE

The IC was fabricated at Taiwan Semiconductor Manufacturing Company's (TSMC) 180-nm CMOS process. TSMC's 180-nm process features a single poly and six metal layers with low-k dielectrics at a supply voltage of 1.8 volts for core logic and 3.3 volts for I/O's. DirAc is packaged in a 276 tape ball grid array (TBGA) measuring 27 mm per side. This custom package provides for 144 signal inputs and outputs, and it includes an internal concentric power ring on the tape to supply power. This allows for over 700 power and ground supply connections to the DirAc die.

The DirAc die, shown in Figure 8, measures 9.40 x 9.36 mm and contains 21.5 million transistors with 1.3 million placeable cells. The prototype fabrication run for the DirAc chip yielded at 82% (148 out of a potential lot size of 180 devices). The chips were functionally tested for all operational modes in November 2002. The total number of test vectors exceeded 150 million test cycles. DirAc's clock frequency exceeded required speed of 40.92 MHz by a factor of two under nominal conditions, providing ample margin. The design operates over the military temperature range of -55° to 125° C. A complete suite of electrical and timing tests showed that the IC met design specifications.




Figure 8. Die Microphotograph of the DirAc IC

Operational Features

The DirAc IC supports either 10 or 5 msec coherent integration in support of different features and modes of M code [11]. DirAc provides a maximum of 128 noncoherent integrations and supports programmable detection processing. Detection reports are available through a simple memory map interface. On-chip 4 Kbits of dual-port RAM provide for storage of two detection reports which include raw peak data, detection threshold levels, calculated noise floor, and time and frequency location.

DirAc supports a power management mode that allows the user to selectively power down individual short-time CMFs as well as the overall device.

The DirAc IC also supports an extensive test/data collection mode. The DirAc datapath has embedded hardware control, snap-shot memory, and dedicated inputs and outputs to support data collection for post-analysis and pre-fabrication verification. Data can be collected at each major processing element in DirAc's datapath through the external noncoherent integration memory by leveraging the hardware test fabric. For example, data can be collected from the short-time CMFs, FFT, or detection processing along with any specified time-frequency tile. Similarly, the test fabric can reduce simulation time. The DirAc requires approximately 1 million clock cycles to process one time-frequency tile at 40.92 MHz. The test structures allow for preloading the CMF bank in a parallel fashion reducing the required simulation cycles.

External Noncoherent Memory Interface

An external memory provides intermediate storage of the CAF during noncoherent integration. At each clock cycle, a column of the time-frequency time is read from memory, modified by DirAc, and written back to external memory. The number of memory bits required is as follows

$$MemorySize = FrequencyBins \times Timeoffsets \times \frac{Bits}{Sample}.$$

For 10 and 5 msec coherent integration modes this results in 14 and 7 Mbit respectively. The memory bandwidth required is

$$MemoryBandwidth = 2 \times FrequencyBins \times Fs \times \frac{Bits}{Sample}.$$

This results in 2.94 Gbit/sec for 10 or 5 msec coherent integration modes.

The external memory selected is a 512K x 36 Quad Data Rate (QDR) memory from Micron Technologies Incorporated, which supports two writes and two reads per clock cycle. The memory device measures 15 x 13 mm. The DirAc interface is operated at 40.92 MHz.

Power and Energy

The DirAc IC dissipates 1.9 Watts when performing 10 msec coherent integrations and 1.1 Watts for 5 msec coherent integrations. DirAc also provides a power management mode for selectively clock-gating the code matched filter bank. In standby mode the DirAc IC dissipates 10 mW.

To assess average power consumption in an operational context, suppose a handheld receiver must support a 72 hour mission. Most of the time, the receiver is in timekeeping mode, but the user needs to obtain a fix periodically (once every 15 minutes, or once every 4 hours, or once every 24 hours). In time keeping mode, the receiver uses a 1 part per million timekeeping circuit, and must search ± 400 Hz of frequency uncertainty. The acquisition circuit uses 5 msec coherent integration time and performs 60 noncoherent integrations to obtain adequate detection performance.

The receiver is powered by two AA alkaline batteries that provide 1400 mA-hr at 3.0 volts. If the DC-DC converter has 95% efficiency, the batteries provide 14,364 W-sec of energy.

Under these conditions, the greatest power consumption occurs when the fix is updated every 15 minutes, rather than with longer update intervals. DirAc IC has to operate for 0.3 seconds each time the receiver is activated. Over the 72 hours, approximately 100 W-sec of energy is used, or an average of less than 0.4 mW over the mission. This represents only 0.7% of the total available battery energy provided by the AA-alkaline batteries.

Acquisition Performance

Acquisition performance of the DirAc IC can be predicted using standard theory. The output signal-to-noise-plus interference ratio (SNIR) after a crosscorrelation is given by

$$\rho_o = \frac{T}{L} \frac{0.25C}{N_0 + J_0} \quad (1)$$

where T is the coherent integration time used in the correlations, L is the implementation loss expressed as a number greater than or equal to unity, C is the received signal power, the factor of 0.25 accounts for splitting the received signal power into four distinct segments (upper and lower sidebands, even and odd spreading symbols) in each coherent integration time, N_0 is the power spectral density of the thermal noise at the receiver front end, and J_0 is the effective power spectral density of the received jamming signal. Details on typical implementation losses can be found in [3, 14].

Thus, for a given received signal power, coherent integration time, implementation loss, receiver noise level, and effective jamming level, (1) yields the output SNIR.

The detection probability is found using the generalized Marcum Q function. Using the notation $P_N(X, Y)$ [15] as the probability that the random variable with $2N$ degrees of freedom and SNIR of X exceeds threshold value of Y , allows the detection probability to be expressed

$$P_d = P_{4N_n}(\rho_o, T) \quad (2)$$

where N_n is the number of coherent integrations times used, and T is the detection threshold calculated to provide the needed false alarm probability for the given number of noncoherent integrations. The factor of four in the subscript in (2) accounts for the fact that the number of complex quantities being noncoherently combined is four times the number of coherent integration times used, reflecting the combination of upper and lower sidebands, and even and odd spreading symbols).

The expressions (1) and (2) can be used to determine the number of coherent integration times needed to achieve a specific detection probability at a given false alarm probability.

The time (in seconds) to search the initial time uncertainty of $\pm \Delta$ sec and an initial frequency uncertainty of $\pm \Phi$ Hz is then

$$T_{search} = N_n \left\lceil \frac{\Delta}{T} \right\rceil \left\lceil \frac{\Phi T}{N_{STC}} \right\rceil \quad (3)$$

where T is the coherent integration time, N_{STC} is the number of short-time correlations within the coherent integration time, and $\lceil x \rceil$ is the smallest integer greater than x .

FUTURE DIRECTIONS

The DirAc IC prototype provides a baseline for validating the practicality of using direct acquisition to acquire the M-code signal. Next generation acquisition circuits can use more aggressive architectures and can leverage more advanced silicon process technologies to obtain even lower-power, lower-cost, and smaller-size acquisition circuits with even higher performance. Currently, 130 nm CMOS technology is widely available and 90 nm CMOS technology is emerging. Migration to the next generation 130 nm process alone will reduce power dissipation by a factor of three while reducing die area and cost of nonrecurring engineering by a factor of two, compared to the prototype DirAc IC. Migration to 90 nm will produce similar improvements yet again over the 130 nm implementation.

Since the DirAc IC design began in 2001, we have considered various enhancements motivated by a combination of lessons learned, refinements of the M-code signal design, evolution of operational concepts for the M-code signal, and developments in semiconductor technology. The enhancements described in this section are divided into two groups: those with minor effect on transistor count and clock speed, and those with more significant effect on transistor count and clock speed.

The following enhancements would affect transistor count and clock speed of prototype DirAc by minimal amounts.

- Minor changes to design logic that would support 256 noncoherent integrations when 5 msec coherent integration time is used.
- Adding another 32-point FFT to allow parallel computation of two CAFs when 5 msec coherent integration time is used.
- Increasing the FFT size to 64 points, providing either lower implementation loss with the same frequency coverage, or twice the frequency coverage with the same implementation loss. If the larger FFT is used to extend the frequency coverage, the size of the external memory used for noncoherent integration would be doubled.
- Postprocessing of magnitude CAFs to reduce worst-case implementation losses.

These enhancements could further improve performance by more than a factor of two. Since they have little effect on complexity of the IC, implementation in 130 nm or 90 nm technology provides a small IC with very low power consumption and excellent capability.

More significant enhancements to the prototype DirAc are also under consideration, including the following:

- The coherent integration time could be doubled to 20 msec, providing significant performance benefits at low data rate.
- The sampling rate could also be doubled with the coherent integration time kept the same as in the prototype, providing lower implementation loss and thus better performance.

While these variants would approximately double the number of transistors relative to the Version 1 prototype, the use of 130 nm or 90 nm technology would make the resulting IC smaller and lower power than the DirAc prototype is today.

SUMMARY

This paper has described the DirAc prototype integrated circuit for direct acquisition of the M code signal. DirAc represents the beginning of a new generation of direct acquisition circuitry, where sophisticated processing algorithms and advanced architectures combine with tens of thousands of physical correlators to execute millions of virtual correlations in parallel. DirAc demonstrates that this level of capability can be obtained using mature semiconductor technology, yielding circuits that are high-yield, small, and low-power.

DirAc's code-matched filter architecture allows for low clock rates and minimizes on-chip storage. The resulting architecture is dominated by systolic arrays that can be designed and laid out for maximum efficiency, then replicated. The systolic array architecture also facilitates clock and power distribution.

The IC has been extensively tested in an IC tester, demonstrating that its operation matches design specifications and computer simulations. The IC is currently being integrated into a test receiver to allow more extensive testing.

As semiconductor technology advances, enabling the use of even more transistors and of higher clock rates, there are many opportunities to develop even more capable circuits for direct acquisition of the M-code signal, either using extensions of the DirAc architecture or entirely different approaches.

ACKNOWLEDGMENT

This work was supported by the US Air Force under contract FA8721-04-C-0001.

REFERENCES

1. B. C. Barker, et al., "Overview of the GPS M-code Signal," *Proceedings of The Institute of Navigation National Technical Meeting*, Anaheim, California, January 2000.
2. J. W. Betz, "Binary Offset Carrier Modulations for Radionavigation," *NAVIGATION: Journal of The Institute of Navigation*, Vol. 48, No. 4, Winter 2001-2002.
3. R. Wolfert, et al., "Direct P(Y)-Code Acquisition under a Jamming Environment," *IEEE 1998 Position Location and Navigation Symposium*, April 1998, pp. 228-235.
4. L. E. Franks, *Signal Theory*, Dowden & Culver, 1981.
5. P. Fishman and J. W. Betz, "Predicting Performance of Direct Acquisition for the M Code Signal," *Proceedings of ION 2000 National Technical Meeting*, Institute of Navigation, January 2000.
6. G. R. Povey, et al., "Simplified Matched Filter Receiver Designs for Spread Spectrum Communications Applications," *Electronic and Communication Engineering Journal*, April 1993.
7. M. K. Sust, et al., "Rapid Acquisition Concept for Voice Activated CDMA Communications," *IEEE Globecom 90*, December 1990.
8. R. G. Davenport, "FFT Processing of Direct Sequence Spreading Codes Using Modern DSP Microprocessors," *Proceedings of the IEEE National Aerospace and Electronics Conference, NAECON 1991*, Vol. 1, pp. 98 -105.
9. W. Jones and J. Luecke, "Performance of a Parallel Time-Frequency Direct Sequence Spread-Spectrum Acquisition Scheme," *IEEE Military Communications Conference 1990, MILCOM '90, Conference Record*, Vol. 1, pp.5.3.1-5.3.5.
10. C. Spillard, et al., "A Serial-Parallel FFT Correlator for PN Code Acquisition from LEO Satellites," *IEEE SSSTA 1998*, South Africa, September 2-4, 1998, pp. 446-448.
11. P. Dafesh, et al., "Description and Analysis of Time-Multiplexed M-Code Data," *Proceedings of ION 58th Annual Meeting/CIGTF 21st Guidance Test Symposium*, June 2002, Albuquerque, NM.
12. J. W. Betz, "Effects of Uncompensated Relative Time Companding on a Broadband Cross Correlator," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, Vol. ASSP-33, No. 3, June 1985, pp. 505-510.
13. J. W. Betz, "Performance of the Deskewed Short-Time Correlator," in *Coherence and Time Delay Estimation*, Edited by G. Clifford Carter, IEEE Press, 1993.
14. R. C. Dixon, *Spread Spectrum Systems*, Wiley-Interscience, New York, 1984.
15. D. A. Shnidman, "The Calculation of the Probability of Detection and the Generalized Marcum Q-Function," *IEEE Transactions on Information Theory*, Vol. 35, No 2, March 1989.